




Doc Code: AP.PRE.REQ

PTO/SB/33 (07-05)

Approved for use through xx/xx/200x. OMB 0651-00xx


U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

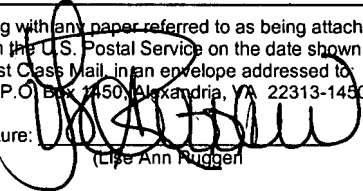
PRE-APPEAL BRIEF REQUEST FOR REVIEW		Docket Number (Optional) EQLC-P01-003	
	Application Number 10/762,984	Filed January 21, 2004	
	First Named Inventor Koning et al.		
	Art Unit 2188	Examiner D. T. Doan	
<p>Applicant requests review of the final rejection in the above-identified application. No amendments are being filed with this request.</p> <p>This request is being filed with a notice of appeal.</p> <p>The review is requested for the reason(s) stated on the attached sheet(s). Note: No more than five (5) pages may be provided.</p> <p>I am the</p> <p><input type="checkbox"/> applicant /inventor.</p> <p><input type="checkbox"/> assignee of record of the entire interest. See 37 CFR 3.71. Statement under 37 CFR 3.73(b) is enclosed. (Form PTO/SB/96)</p> <p><input checked="" type="checkbox"/> attorney or agent of record. Registration number <u>46,698</u></p> <p><input type="checkbox"/> attorney or agent acting under 37 CFR 1.34. Registration number if acting under 37 CFR 1.34. _____</p> <p> _____ Signature Richard M. Feustel, Esq. _____ Typed or printed name (617) 951-7760 _____ Telephone number September 10, 2007 _____ Date</p> <p>NOTE: Signatures of all the inventors or assignees of record of the entire interest or their representative(s) are required. Submit multiple forms if more than one signature is required, see below*.</p> <p><input checked="" type="checkbox"/> *Total of <u>1</u> forms are submitted.</p>			

I hereby certify that this paper (along with any paper referred to as being attached or enclosed) is being deposited with the U.S. Postal Service on the date shown below with sufficient postage as First Class Mail, in an envelope addressed to: MS AF, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Dated: September 10, 2007Signature:  (Lisa Ann Ruggeri)



I hereby certify that this paper (along with any paper referred to as being attached or enclosed) is being deposited with the U.S. Postal Service on the date shown below with sufficient postage as First Class Mail in an envelope addressed to: MS AF, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Dated: September 10, 2007 Signature: 

(Lisa Ann Ruggier)

Docket No.: EQLC-P01-003
(PATENT)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:
Koning et al.

Application No.: 10/762,984

Confirmation No.: 5998

Filed: January 21, 2004

Art Unit: 2188

For: BLOCK DATA MIGRATION

Examiner: D. T. Doan

PRE-APPEAL BRIEF REASONS FOR REVIEW

MS AF
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

In response to the Final Office Action mailed March 21, 2007, Applicants request a panel review of the pending rejections prior to proceeding with the full appeals process. Applicants enclose with these remarks the requisite Notice of Appeal and extension of time.

The Final Office Action is the third Office Action issued in this case. Applicants have amended independent claims 1 and 14 twice during prosecution, arguing patentability based on the amendments made. In view of the fact that the current rejection is based on the combination of four disparate and unrelated references, Applicants believe it is appropriate to seek a panel review and to present additional arguments concerning various deficiencies in the rejection.

The Claims. Applicants' claims are directed to approaches for resource migration in a server-based partitioned-resource storage system. In such a system, a plurality of storage *servers* has "a set of resources partitioned thereon."¹ That is, portions of a given resource are stored across

¹ In some embodiments, a partitioned resource storage system may also store unpartitioned resources, *i.e.*, a server in the system may be responsible for storing an entire resource (*see, e.g.*, ¶ 46).

various storage servers in the system (*see, e.g.*, FIG. 4 and ¶¶ 44 et seq.).² The storage servers have respective load monitor processes that communicate with each other to determine a measure of loading for each respective storage server. Based on the determined measure of loading, resources are transferred from one server to another.

A write detect process detects when a resource is in the process of being moved between first and second servers, and writes copies of the resource to both of the servers. A write failure on the target server (*i.e.*, the claimed “second server”) is handled in a simple and efficient way – by restarting the migration process for the resource having the write request issued to it. In this way, the simultaneous write to the source server (*i.e.*, the claimed “first server”) is leveraged. The affected resource is simply re-migrated, “ensur[ing] that the write request is propagated to the second server.”

Mashayekhi. The primary reference used in the rejection is Mashayekhi et al. US 2003/0074596. Applicants respectfully submit that Mashayekhi is inapposite to the claims because, in addition to missing the “claim’s detail of the server’s processes” (Office Action p. 3) as the Examiner acknowledges, Mashayekhi does not disclose the claimed servers/partitioned resources. The Examiner cites elements 15 and 16 of FIG. 1 for the claimed servers/partitioned resources, but this is clearly an error. Mashayekhi is a *clustered* system, not a server-based partitioned system as claimed (Mashayekhi ¶¶ 3-7). As explained in Mashayekhi, a clustered system includes a shared external resource used by all of the servers in the system (Mashayekhi ¶ 3). Resources are not partitioned across storage servers. Applicants believe that the Examiner is perhaps somehow trying to suggest that the cluster metadata is a partitioned resource, but this too is clear error.³ The metadata is not partitioned across Mashayekhi’s Network Interface Cards (NICs), but rather each NIC stores an entire copy of the metadata (Mashayekhi ¶ 8; Abstract). Moreover, there is no reason, other than an impermissible amount of hindsight, one would completely redesign Mashayekhi to have the claimed storage servers/partitioned resources. Mashayekhi’s invention is directed to a complex communications scheme for ensuring consistent metadata across servers using

² Each storage server has its own associated storage device or devices.

³ The metadata is not part of the data stored on storage 15. As defined in Mashayekhi, the metadata of a cluster (as opposed to metadata generally) “includes information describing the logical relationship of servers in the cluster and their association with the services provided by the cluster.” Mashayekhi ¶ 3.

custom NICs in a distinctly different clustered system. This reference and Applicants' claims are ships passing in the night.

Umberger. The Examiner next relies on Umberger et al. US 6,957,433 to provide the "claim's detail of the server's processes." But the Examiner's analysis ignores features of the claims. The claims state that each of the storage servers has a respective load monitor process capable of communicating with other load monitor processes for other respective servers for generating a measure of loading "on respective ones" of the storage servers. Applicants have searched Umberger and cannot find any such "respective" load monitoring processes for storage servers. Indeed, any such analysis is noticeably absent from the Examiner's analysis on p. 3 of the Office Action. The Examiner cites the servers of FIG. 5, but Umberger clearly states that there is a single "management computer 504" for balancing load (Umberger 8:37-57). This centralized approach for determining load is not what is claimed. All storage systems 502 communicate with the server 504, not with other storage systems. Server 504 is plainly not a storage system, it is a "management computer." The Examiner also cites FIG. 3. Again, Applicants cannot find any disclosure where any components of FIG. 3 are on multiple respective storage servers and communicate for purposes of determining a measure of load for each respective server. Umberger discloses instead centralized approaches. For example, column 6, lines 56-67 disclose one device for performing functions of 302-305 (e.g., "*a* service component of data processing system [singular], such as, for instance, as processor core, input/output device, communications port, or storage system [all singular]" (6:62-64 (emphasis and comments added))). Thus, the combination so far would not result in the claimed load monitor process (and the other aforementioned missing elements).

Blumenau. The third reference in the combination in Blumenau US 2004/0080558. As an initial matter, Blumenau does not make up for any of the aforementioned deficiencies of the other references in this combination. The Examiner relies upon Blumenau for providing a simultaneous write process, but completely disregards Blumenau's solution for handling write failures during such a process.⁴ What motivation could there be, to one of ordinary skill in the art, to completely

⁴ Blumenau does not disclose such an approach for handling target-server write failures. It instead discloses a much more involved and complicated recovery process that compares the source and target data with state information (see paragraph 48, lines 2-5 and paragraph 49, lines 2-7). The state information includes information that tracks the state of

redesign Blumenau and arrive at Applicants' approach? The rejection states the answer is found in Aditya et al. US 6,122,681.

Aditya. Applicants' claims recite that, in response to a write failure, the write-detect process "restarts the migration process for the resource to ensure that the write request is propagated to the second server." The Examiner relies upon Aditya's column 2, lines 7-26 for this feature. The pertinent portion is provided here for the Panel's convenience:

The physical layer is the hardware which in a network includes the network controller and the physical link. If the physical link is linear, such as Ethernet, a carrier sense multiple access/collision detection (CSMA/CD) system is used in which a node sends a signal that every other node detects but only the addressed node interprets as useful data. If two nodes send signals at the same time, a collision occurs and both backoff, wait for a unique random amount of time and then try again.

Applicants respectfully submit that CSMA/CD is so far afield from anything having to do with Applicants' invention, that the fact the Examiner had to go this far to find something to map to Applicants' claimed feature actually demonstrates the patentability of Applicants' claims. CSMA/CD is a well-known physical layer protocol. One of ordinary skill would *never* be motivated by such a disclosure to completely redesign Blumenau and arrive at Applicants' application-layer approach for solving write failures during a migration in a storage server based partitioned storage system. More accurately, if one of ordinary skill were to use the teachings of Aditya, all he or she would produce is a storage system that employs CSMA/CD at the physical layer. There is no basis for concluding otherwise.

Conclusion. Applicants have shown numerous deficiencies in the four-reference obviousness rejection of the Final Office Action. A *prima facie* case of obviousness has not been made. The rejection of all of the claims should be withdrawn.

each storage element in the system (e.g., a count which indicates a number of data operations performed on a particular storage location) (see paragraph 51) and is stored for use in the recovery process in the event of a write failure. Blumenau's recovery process copies the "good data" from the storage location where the write completed successfully to the other location based on the state information (see paragraph 48, lines 5-7 and paragraph 85, lines 9-12). The data in the location where the most recent write occurred is relied upon as "good data," based on the state information (see paragraph 48, lines 13-17 and paragraph 53, lines 5-12). Alternatively, Blumenau's recovery process invalidates data stored at both locations (see paragraph 48, lines 20-23 and paragraph 85, lines 9-11). This is entirely different from Applicants' simplified approach of restarting the migration process for the resource.

Applicants believe that all fees due in connection with this submission have been appropriately provided. However, if there are any other fees due in connection with the filing of this Response, please charge the fees to Deposit Account No. 18-1945, under Order No. EQLC-P01-003 from which the undersigned is authorized to draw.

Dated: September 10, 2007

Respectfully submitted,

By 

Richard M. Feistel, Jr., Esq.

Registration No.: 46,698

ROPES & GRAY LLP

One International Place

Boston, Massachusetts 02110

(617) 951-7000

(617) 951-7050 (Fax)

Attorneys/Agents For Applicant